

The Stochastic complexity of spin models: How simple are simple spin models?

Alberto Beretta,¹ Claudia Battistin,² Clélia de Mulatier,¹ Iacopo Mastromatteo,³ and Matteo Marsili¹

¹*The Abdus Salam International Centre for Theoretical Physics (ICTP), Strada Costiera 11, I-34014 Trieste, Italy*

²*Kavli Institute for Systems Neuroscience and Centre for Neural Computation, Olav Kyrres gate 9, 7030 Trondheim, Norway*

³*Capital Fund Management, 23 rue de l'Université, 75007 Paris, France*

Simple models, in information theoretic terms, are those with a small stochastic complexity. We study the stochastic complexity of spin models with interactions of arbitrary order. Invariance with respect to bijections within the space of operators allows us to classify models in complexity classes. This invariance also shows that simplicity is not related to the order of the interactions, but rather to their mutual arrangement. Models where statistical dependencies are localized on non-overlapping groups of few variables (and that afford predictions on independencies that are easy to falsify) are simple. On the contrary, fully connected pairwise models, which are often used in statistical learning, are highly complex because of their extended set of interactions.

Science, as the endeavour of reducing complex phenomena to simple models, has been instrumental to solve practical problems. The Big Data deluge has shown that understanding is no more necessary to solve problems [1, 2] such as image or speech recognition and language translation [3]. A statistical model trained on a sufficiently large number of instances of the solution can learn how to generalise from examples in ways that resemble the “unknown rules” that humans follow [4, 5]. These models are *simple* in the sense that they must be easy to evaluate, train and/or to infer. They offer *simple* interpretations in terms of low order (typically pairwise) dependencies, which in turn afford an explicit graph theoretical representation [6]. Their aim is not to uncover fundamental laws but to “generalise well”, i.e. to describe well yet unseen data. For this reason, machine learning relies on “universal” models that are apt to describe any possible data on which they can be trained [7], using suitable “regularisation” schemes in order to tame parameter fluctuations (overfitting) and achieve small generalisation error [8].

Scientific models, instead, are the simplest possible descriptions of experimental results. A physical model is a representation of a real system and its structure reflects the laws and symmetries of Nature. It predicts well not because it generalises well, but rather because it captures essential features of the specific phenomena that it describes. It should depend on few parameters and is designed to provide predictions that are easy to be falsified [9]. For example, Newton’s laws of motion are consistent with momentum conservation, a fact that can be checked in scattering experiments.

The intuitive notion of a “simple model” hints at a succinct description, one that requires few bits [10]. Minimum Description Length (MDL) [11, 12] provides a quantitative framework for “counting complexity” of models in bits. So, what are simple models according to MDL and which intuitive notion of simplicity do they conform to?

This paper addresses this issue in the context of the exponential family of spin models, for which the probability distribution over the n spin variables $\mathbf{s} = (s_1, \dots, s_n)$, taking values $s_i = \pm 1$, is given by:

$$P(\mathbf{s} | \mathbf{g}, \mathcal{M}) = \frac{e^{\sum_{\mu \in \mathcal{M}} g^\mu \phi^\mu(\mathbf{s})}}{Z_{\mathcal{M}}(\mathbf{g})}, \quad (1)$$

$$\text{where } Z_{\mathcal{M}}(\mathbf{g}) = \sum_{\mathbf{s}} e^{\sum_{\mu \in \mathcal{M}} g^\mu \phi^\mu(\mathbf{s})}. \quad (2)$$

Here, $\{\phi^\mu(\mathbf{s}), \mu \in \mathcal{M}\}$ is a collection of (monomial) spin operators, $\phi^\mu(\mathbf{s}) = \prod_{i \in \mu} s_i$ where μ is a subset of the n spins, that characterises the model \mathcal{M} (see [13] Sec. 0). The g^μ are the conjugate parameters [14] and the normalisation $Z_{\mathcal{M}}(\mathbf{g})$ is the partition function.

For this class of models, the complexity can be defined unambiguously within MDL as the number of bits needed to specify *a priori* the parameters $\hat{\mathbf{g}}$ that best describe a dataset $\hat{\mathbf{s}} = (\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(N)})$ consisting of N samples independently drawn from the distribution $P(\mathbf{s} | \mathbf{g}, \mathcal{M})$ for some unknown \mathbf{g} (see [13] Sec. 0). Asymptotically for $N \rightarrow \infty$, the MDL complexity is given by [15, 16]:

$$\log \sum_{\hat{\mathbf{s}}} P(\hat{\mathbf{s}} | \hat{\mathbf{g}}, \mathcal{M}) \simeq \frac{|\mathcal{M}|}{2} \log \left(\frac{N}{2\pi} \right) + c_{\mathcal{M}}. \quad (3)$$

The first term, which is the basis of the Bayesian Information Criterion (BIC) [17, 18], captures the increase of the complexity with the number $|\mathcal{M}|$ of model’s parameters and with the number N of data points. This accounts for the fact that the uncertainty in each parameter $\hat{\mathbf{g}}$ decreases with N as $N^{-1/2}$, so its description requires $\sim \frac{1}{2} \log N$ bits. The second term is the *stochastic complexity* [19]:

$$c_{\mathcal{M}} = \log \int d\mathbf{g} \sqrt{\det \mathbb{J}(\mathbf{g})}, \quad (4)$$

where $\mathbb{J}(\mathbf{g})$ is the Fisher Information matrix with entries

$$J_{\mu\nu}(\mathbf{g}) = \frac{\partial^2}{\partial g^\mu \partial g^\nu} \log Z_{\mathcal{M}}(\mathbf{g}). \quad (5)$$

$c_{\mathcal{M}}$ encodes the intrinsic notion of simplicity we are interested in. Its expression (4) coincides with the definition of *geometric complexity* (the penalty that should be applied to the model \mathcal{M} because of its geometry) in the context of Bayesian model selection when using Jeffreys' prior over the parameters \mathbf{g} [13, 18, 20, 21]. Our aim is to understand how the stochastic complexity depends on the structure of the model \mathcal{M} .

The class of models we study contains pairwise interaction models, where the operators $\phi^\mu(\mathbf{s})$ are single spins s_i or product of two spins $s_i s_j$, for $i, j \in \{1, \dots, n\}$. There has been a surge of recent interest in the inference of spin models [22] from high dimensional data, most of which was limited to pairwise models. This is partly because, for large n , inference is computationally challenging already in this case, but more importantly because, since the number of k -spin interactions grows as n^k , data size is hardly sufficient to go beyond $k = 2$. Ref. [23] recently suggested that even for data generated by models with higher order interactions, pairwise models may provide a *sufficiently* accurate description of the data. Within the class of pairwise models, L1 regularisation [24] has proven to be a remarkably efficient heuristic of model selection (but see also [25]). We remark that the models of Eq. (1) can be derived as the maximum entropy distributions that are consistent with the requirement that the model reproduces the empirical averages of the operators $\phi^\mu(\mathbf{s})$ for all $\mu \in \mathcal{M}$ on a given dataset [26, 27]. In other words, empirical averages of $\phi^\mu(\mathbf{s})$ are sufficient statistics, i.e. their values are enough to compute the maximum likelihood parameters $\hat{\mathbf{g}}$. Therefore the choice of the operators ϕ^μ in \mathcal{M} inherently entails a sharp separation between relevant variables (the sufficient statistics) and irrelevant ones, which might not be justified *a priori*. For example, if statistical inference assumes pairwise interactions, it might be blind to relevant patterns in the data resulting from higher order interactions.

In this light, the study of exponential models with interactions of arbitrary order sheds light on which models should be privileged within a fully Bayesian approach and, ultimately, on which operators are relevant to describe a dataset for a given N . In particular, we want to understand whether pairwise models have distinctive virtues that make them particularly efficient in statistical learning (such as the *pairwise sufficiency* of Ref. [23]).

Let's start by showing that low order interactions do not have a privileged status and are not necessarily related to low complexity $c_{\mathcal{M}}$, with the following argument: Alice is interested in finding which model \mathcal{M} best describes a dataset $\hat{\mathbf{s}}$. Bob is interested in the same problem, but his dataset $\hat{\boldsymbol{\sigma}}$ is related to Alice's dataset by a *gauge transformation*. The latter is defined as a bijective transformation between the n random variables \mathbf{s} of Alice and those of Bob, $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$, that corresponds to a bijection of the set of operators $\phi^\mu(\mathbf{s}) \rightarrow \phi^\mu(\boldsymbol{\sigma})$ (see the examples in Fig. 1 and [13] Sec. 1). This induces a bi-




FIG. 1. Example of gauge transformations between models with $n = 4$ spins. Models are represented by diagrams: spins are full dots \bullet in presence of a local field, or empty dots \circ otherwise; blue lines are pairwise interactions ($\phi^\mu = s_i s_j$); orange triangles denote 3-spin interactions ($\phi^\mu = s_i s_j s_k$); and the 4-spin interaction ($s_1 s_2 s_3 s_4$) is a filled blue triangle. Note that model a) has all its interactions grouped on 3 spins; the gauge transformations leading to this model are shown along the arrows. All the models belong to the same complexity class, with $|\mathcal{M}| = 7$, $\lambda = 4$ and a number of independent operators $n_{\mathcal{M}} = 3$ (e.g. s_1, s_2 and s_3 in model a) – see tables in [13] Sec. 6). The class contains in total 15 models, which are grouped, with respect to the permutation of the spins, behind the 4 representatives shown here with their multiplicity ($\times m$).

jective transformation between Alice's models and those of Bob, as shown in Fig. 1, that preserves the number of interactions $|\mathcal{M}|$. Whatever conclusion Bob draws on the relative likelihood of models can be translated into Alice's world, where it has to coincide with Alice's result. It follows that two models \mathcal{M} and \mathcal{M}' related by a gauge transformation must also have the same complexity $c_{\mathcal{M}} = c_{\mathcal{M}'}$. In particular, pairwise interactions can be mapped to interactions of any order (see Fig. 1), and, consequently, low order interactions are not necessarily simpler than higher order ones.

Observe that models connected by gauge transformations have remarkably different structures. In Fig. 1, model a) has all the possible interactions concentrated on 3 spins, having the properties of a *simplicial complex* [28]; however, its gauge-transformed counterparties are not simplicial complexes. Model d) is invariant under any permutations of the four spins, whereas the other models have a lower degree of symmetry under permutations (see the different multiplicities in Fig. 1).

Gauge transformations are discussed in more details in [13] Sec. 1. Here let it suffice to say that the number of such bijections is:

$$\mathcal{N}_{GT}(n) = 2^{n^2} \prod_{k=1}^n (1 - 2^{-k}), \quad (6)$$

which is much less than the number $2^n!$ of possible bijections of the set of 2^n states into itself.

Gauge transformations allow us to divide the set of all models into equivalence classes, which we call *complexity classes*. Models belonging to the same class are related to each other by a gauge transformation (that is the equivalence relation), and thus have the same complexity $c_{\mathcal{M}}$. This classification suggests the presence of “quantum numbers” (invariants) in terms of which models can be classified. These invariants emerge explicitly when writing the cluster expansion of the partition function [29–31] (see [13] Sec. 2):

$$Z_{\mathcal{M}}(\mathbf{g}) = 2^n \left(\prod_{\mu \in \mathcal{M}} \cosh(g^\mu) \right) \sum_{\ell \in \mathcal{L}} \prod_{\mu \in \ell} \tanh(g^\mu). \quad (7)$$

The sum runs on the set \mathcal{L} of all possible *loops* ℓ that can be formed with the operators $\mu \in \mathcal{M}$. A loop is any subset $\ell \subseteq \mathcal{M}$ of operators such that each spin s_i occurs zero or an even number of times in the product $\prod_{\mu \in \ell} \phi^\mu(\mathbf{s})$, which is then equal to 1 for any value of the variables \mathbf{s} . The set \mathcal{L} of loops of any model \mathcal{M} has the structure of a finite group: if $\ell_1, \ell_2 \in \mathcal{L}$, then $\ell_1 \oplus \ell_2$ is also a loop of \mathcal{M} , where \oplus is the XOR operator on spins (see [13] Sec. 3); this structure is an invariant of the class. It follows that each loop of \mathcal{L} can be generated by a *minimal generating set* of λ loops, and that each model has $|\mathcal{L}| = 2^\lambda$ loops (including the empty loop $\{\emptyset\}$). We found that λ is related to the number $|\mathcal{M}|$ of operators of the model by $\lambda = |\mathcal{M}| - n_{\mathcal{M}}$, where $n_{\mathcal{M}}$ is the number of *independent operators* of \mathcal{M} , i.e. the maximal number of operators that can be taken in \mathcal{M} without forming any loop. This implies that λ attains its minimal value, $\lambda = 0$, for models with only independent operators, and its maximal value, $\lambda = 2^n - 1 - n$, for the *complete model* $\overline{\mathcal{M}}$, that contains all the $|\overline{\mathcal{M}}| = 2^n - 1$ possible operators among n spins. Finally, gauge transformations also imply a duality relation, by which to each class of models with $|\mathcal{M}|$ operators corresponds a class of models with the $2^n - 1 - |\mathcal{M}|$ complementary operators (see [13] Sec. 3).

Coming to the quantitative estimate of the complexity, $c_{\mathcal{M}}$ generally depends on the extent to which ensemble averages of the operators $\phi^\mu(\mathbf{s})$ in the model $\mu \in \mathcal{M}$ constrain each other. This appears explicitly by rewriting Eq. (4) as an integral over the ensemble averages of the operators, $\boldsymbol{\varphi} = \{\langle \phi^\mu \rangle, \mu \in \mathcal{M}\}$, exploiting the bijection between the parameters \mathbf{g} and their dual parameters $\boldsymbol{\varphi}$ and re-parameterization invariance [21, 32]:

$$c_{\mathcal{M}} = \log \int_{\mathcal{F}} d\boldsymbol{\varphi} \sqrt{\det \mathbb{J}(\boldsymbol{\varphi})}, \quad (8)$$

where $\mathbb{J}(\boldsymbol{\varphi})$ is the Fisher Information Matrix in the $\boldsymbol{\varphi}$ -coordinates. Here, the domain \mathcal{F} of integration is over the values of $\boldsymbol{\varphi}$ that can be realised in any empirical sample drawn from the model \mathcal{M} (known in this context as *marginal polytope* [33]) and is related to the mutual constraints between the ensemble averages φ^μ




FIG. 2. Complexity c_k of models with a single loop of length k and $n - k + 1$ free operators, i.e. that are not involved in any loop. Consider for instance a model with n fields (one on each spin) and one $(k - 1)$ -body interaction, or a model formed with a closed chain of k pairwise interactions and $n - k + 1$ free fields. The complexity is $c_k = a_k \log \pi + (n - k + 1) \log \pi$, where $(n - k + 1) \log \pi$ is the contribution of the free operators and $a_k \log \pi$ is the one of the loop. The complexity $c_3 = n \log \pi$ corresponds to the case $k = 3$ ($a_k = 2$, see [13] Sec. 5). Note that this curve is independent of the value of n , as $(c_k - c_3) / \log \pi = a_k - k + 1$. Values of c_k are averaged over 10^3 numerical estimates of the integral in Eq. (4), using 10^6 MC samples each (see [13] Sec. 6). Error bars correspond to the standard deviation of these estimates.

(see [13] Sec. 4 for more details). If the model contains no loop, $\mathcal{L} \setminus \{\emptyset\} = \emptyset$, then $J_{\mu\nu}(\boldsymbol{\varphi}) = [1 - (\varphi^\mu)^2]^{-1} \delta_{\mu\nu}$ is diagonal: the integral in Eq. (8) factorizes and gives $c_{\mathcal{M}} = |\mathcal{M}| \log \pi$. In this case, the variables φ^μ are not constrained at all and the domain of integration is $\mathcal{F} = [-1, 1]^{|\mathcal{M}|}$. If instead the model contains loops, the variables φ^μ become constrained and the marginal polytope \mathcal{F} is reduced. For example, for a model with a single loop of length three (e.g. $\phi^1 = s_1$, $\phi^2 = s_2$ and $\phi^3 = s_1 s_2$), the values of $\boldsymbol{\varphi}$ in $[-1, 1]^3$ are not all attainable, indeed $\mathcal{F} = \{\boldsymbol{\varphi} \in [-1, 1]^3 : |\varphi^1 + \varphi^2| - 1 \leq \varphi^3 \leq 1 - |\varphi^1 - \varphi^2|\}$ is reduced, which decreases the complexity. The complexity for a model with a fixed number $|\mathcal{M}|$ of parameters and a single loop of length k can be computed (see [13] Sec. 6), and is shown in Fig. 2. The complexity c_k increases with the length k of the loop, and saturates at $c_3 + \log \pi = |\mathcal{M}| \log \pi$, which is the value one would expect if the k operators were unconstrained. This is consistent with the idea that longer loops induce weaker constraints among the operators.

The single loop calculation allows computing the complexity of models with non-overlapping loops ($\ell \cap \ell' = \emptyset$ for all $\ell, \ell' \in \mathcal{L}$), for which $c_{\mathcal{M}} = \sum_{\ell \in \mathcal{L}} c_{|\ell|}$. In the general case of models with a more complex loop structure, the explicit calculation of $c_{\mathcal{M}}$ is non-trivial. Yet, the argument above suggests that the complexity $c_{\mathcal{M}}$, for a given number of parameters $|\mathcal{M}|$, should increase with the number $n_{\mathcal{M}}$ of independent operators. Fig. 3 sum-




FIG. 3. Complexity of models for $n = 4$ as a function of the number $|\mathcal{M}|$ of operators: each triangle represents a class of complexity, which contains one or more models (see [13] Sec. 6). For each class, the value of the $c_{\mathcal{M}}$ was obtained from a *representative* of the class; some of them are shown here with their corresponding diagram (same notations as in Fig. 1). The colors of the triangles indicate the values of $n_{\mathcal{M}}$: violet for $n_{\mathcal{M}} = 4$, green for 3, yellow for 2, pink for 1 and red for 0 (model with no operator). Models on the black line are models with only independent operators ($|\mathcal{M}| = n_{\mathcal{M}}$), whose complexity is $c_{\mathcal{M}} = |\mathcal{M}| \log \pi$; models on the dashed curve are complete models, whose complexity is given in Eq. (9). Complexity classes with the same values of $|\mathcal{M}|$ and $n_{\mathcal{M}}$ have the same value of $\lambda = |\mathcal{M}| - n_{\mathcal{M}}$, i.e. the same number of loops $|\mathcal{L}|$, but with a different structure.

marises the results for all models with $n = 4$ spins (see [13] Sec. 6), and supports this conclusion: for a given value of $|\mathcal{M}|$, classes with lower values of $n_{\mathcal{M}}$ (i.e. with less independent operators) are less complex.

A surprising result of Fig. 3 is that $c_{\mathcal{M}}$ is not monotonic with the number $|\mathcal{M}|$ of operators of the model, increasing first with $|\mathcal{M}|$ and then decreasing. Complete models $\overline{\mathcal{M}}$ turn out to be the simplest (see the dashed curve in Fig. 3). As a consequence, for a given $|\mathcal{M}|$, models that contain a complete model on a subset of spins are generally simpler than models where operators have support on all the spins. For instance, the complexity class displayed in Fig. 1 is the class of models with $|\mathcal{M}| = 7$ operators that has the lowest complexity (see green triangle on the dashed curve in Fig. 3).

Fig. 3 also confirms that pairwise models are not simpler than models with higher order interactions. Indeed, as shown in Fig. 3, the complexity increases when changing the model a) of Fig. 1 into a pairwise models by turning the 3-spin interaction into an external field acting on s_4 . Likewise, the model with all 6 pairwise interactions for $|\mathcal{M}| = 10$ is more complex than the one where one of them is turned into a 3-spin interaction.

It is possible to compute explicitly the complexity of a complete model $\overline{\mathcal{M}}$ with n spins. Indeed, there is a

mapping $g^\mu = 2^{-n} \sum_{\mathbf{s}} \phi^\mu(\mathbf{s}) \log p(\mathbf{s})$ between the $2^n - 1$ parameters g^μ of $\overline{\mathcal{M}}$ and the 2^n probability $p(\mathbf{s})$, also constrained by their normalisation [34]. The complexity in Eq. (4) is invariant under reparametrization [32]. Rewriting this integral in terms of the variables $p(\mathbf{s})$ and using that $\det \mathbb{J}(\mathbf{p}) = \prod_{\mathbf{s}} 1/p(\mathbf{s})$, we find (see [13] Sec. 5):

$$c_{\overline{\mathcal{M}}} = \log \int_0^1 d\mathbf{p} \delta \left(\sum_{\mathbf{s}} p(\mathbf{s}) - 1 \right) \prod_{\mathbf{s}} \frac{1}{\sqrt{p(\mathbf{s})}}, \quad (9)$$

$$= 2^{n-1} \log \pi - \log \Gamma(2^{n-1}). \quad (10)$$

Note that, for $n > 4$, $c_{\overline{\mathcal{M}}}$ becomes negative (for $n = 6$, $c_{\overline{\mathcal{M}}} \simeq -41.5$). This suggests that the class of least complex models with $|\mathcal{M}|$ interactions is the one that contains the model where the maximal number of interactions are concentrated on a small number k of spins (with k the largest integer $k \leq \log_2 |\mathcal{M}|$). This agrees with our calculation for models with a single loop, where the complexity is minimal when the loop is concentrated on the minimal number of operators. On the contrary, models where the interactions are distributed uniformly across the variables (e.g. models with only single spin operators for $n \geq |\mathcal{M}|$ or with non-overlapping sets of loops) have higher complexity.

This leads us to conjecture that stochastic complexity is related to the localisation properties of the set of loops \mathcal{L} (i.e. its group structure) rather than to the order of the interactions: models where the loops $\ell, \ell' \in \mathcal{L}$ have a “large” overlap $\ell \cap \ell'$ are simple, whereas models with an extended homogeneous network of interactions (e.g. fully connected Ising models with up-to pairwise interaction) have many non-overlapping loops $\ell \cap \ell' = \emptyset$ and therefore are rather complex. It is interesting to note that the former (simple models) lend themselves to predictions on the independence of different groups of spins. These predictions suggest “fundamental” properties of the system under study (i.e. invariance properties, spin permutation symmetry breaking), and they are easy to falsify (i.e. it is clear how to devise a statistical test that would test these hypotheses to any given confidence level). On the contrary, complex models (e.g. fully connected pairwise Ising models) are harder to falsify as their parameters can be adjusted to fit reasonably well any sample, irrespectively of the system under study.

The unreasonable effectiveness of mathematical models ultimately resides in finding phenomena that depend on few variables, whose mutual variation is described by *simple models* and is independent of the rest [35]. Simple models are not optimal to generalise, i.e. to describe observed dependencies. Rather, they are designed for spotting independencies that may hint at deeper principles (e.g. symmetries or conservation laws) that may “take us beyond the data” [36]. The observation that evolution and adaptation drive phenotypic variation on low dimensional manifolds [37] and the emergence of sparse

representations from high dimensional data in deep learning [38–40] suggest that simple models will continue playing a role in developing a quantitative understanding of the world around us. We find it remarkable that this notion of simplicity, in the case of spin models, agrees precisely with the information theoretic notion of stochastic complexity. A Bayesian approach to quantitative science should privilege simple models of this type when data are scarce and high dimensional.

ACKNOWLEDGMENTS

C. B. acknowledges financial support from the Kavli Foundation and the Norwegian Research Council’s Centre of Excellence scheme (Centre for Neural Computation, grant number 223262). A. B. acknowledges financial support from International School for Advanced Studies (SISSA).

-
- [1] V. Mayer-Schonberger and K. Cukier, *Big Data: A Revolution That Will Transform How We Live, Work and Think*. (John Murray Publishers, UK, 2013).
 - [2] C. Anderson, Wired Magazine (2008).
 - [3] N. Cristianini, Neural Networks **23**, 466 (2010).
 - [4] Y. LeCun, K. Kavukcuoglu, and C. Farabet, in *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on* (IEEE, 2010) pp. 253–256.
 - [5] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prengel, S. Satheesh, S. Sen Gupta, A. Coates, and A. Ng, ArXiv e-prints (2014), arXiv:1412.5567 [cs.CL].
 - [6] C. Bishop, *Pattern Recognition and Machine Learning*, Information Science and Statistics (Springer, 2006).
 - [7] X. Wu, V. Kumar, J. Ross Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. J. McLachlan, A. Ng, B. Liu, P. S. Yu, Z.-H. Zhou, M. Steinbach, D. J. Hand, and D. Steinberg, Knowledge and Information Systems **14**, 1 (2008).
 - [8] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016) <http://www.deeplearningbook.org>.
 - [9] K. Popper, *The Logic of Scientific Discovery*, Routledge Classics (Taylor & Francis, 2005).
 - [10] N. Chater and P. Vitényi, Trends in cognitive sciences **7**, 19 (2003).
 - [11] J. Rissanen, Automatica **14**, 465 (1978).
 - [12] P. D. Grünwald, *The minimum description length principle* (MIT press, 2007).
 - [13] See Supplemental Material.
 - [14] There is a broader class of models, where subsets $\mathcal{V} \subseteq \mathcal{M}$ of operators have the same parameter, i.e. $g^\mu = g^\nu$ for all $\mu \in \mathcal{V}$. These *degenerate models* are rarely considered in the inference literature, so we confine our discussion to non-degenerate models and refer the reader to [13] Sec. 7 for more discussion.
 - [15] J. J. Rissanen, IEEE Transactions on Information Theory **42**, 40 (1996).
 - [16] J. Rissanen, IEEE Transactions on Information Theory **47**, 1712 (2001).
 - [17] G. Schwarz, The annals of statistics **6**, 461 (1978).
 - [18] I. J. Myung, V. Balasubramanian, and M. A. Pitt, Proceedings of the National Academy of Sciences **97**, 11170 (2000).
 - [19] J. Rissanen, Journal of Computer and System Sciences **55**, 89 (1997).
 - [20] H. Jeffreys, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences **186**, 453 (1946).
 - [21] S. Amari, *Information Geometry and Its Applications*, Applied Mathematical Sciences (Springer Japan, 2016).
 - [22] H. Chau Nguyen, R. Zecchina, and J. Berg, ArXiv e-prints (2017), arXiv:1702.01522 [cond-mat.dis-nn].
 - [23] L. Merchan and I. Nemenman, Journal of Statistical Physics **162**, 1294 (2016).
 - [24] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty, The Annals of Statistics **38**, 1287 (2010).
 - [25] N. Bulso, M. Marsili, and Y. Roudi, Journal of Statistical Mechanics: Theory and Experiment **2016**, 093404 (2016).
 - [26] E. T. Jaynes, Physical Review **106**, 620 (1957).
 - [27] Y. Tikochinsky, N. Z. Tishby, and R. D. Levine, Physical Review A **30**, 2638 (1984).
 - [28] A *simplicial complex*, in our notation, is a model such that, for any interaction $\mu \in \mathcal{M}$, any interaction that involves any subset $\nu \subseteq \mu$ of spins present in $\phi^\mu(\mathbf{s})$ is also contained in the model (i.e. $\nu \in \mathcal{M}$). Ref. [41] discusses exponential models in the context of simplicial complexes.
 - [29] L. Landau and E. Lifshitz, *Statistical Physics*, v. 5 (Elsevier Science, 2013).
 - [30] H. A. Kramers and G. H. Wannier, Physical Review **60**, 263 (1941).
 - [31] A. Pelizzola, Journal of Physics A: Mathematical and General **38**, R309 (2005).
 - [32] S. Amari and H. Nagaoka, *Methods of Information Geometry*, Translations of mathematical monographs (American Mathematical Society, 2007).
 - [33] M. J. Wainwright and M. I. Jordan, Foundations and Trends® in Machine Learning **1**, 1 (2008).
 - [34] I. Mastromatteo, ArXiv e-prints (2013), arXiv:1311.0190 [cond-mat.stat-mech].
 - [35] E. P. Wigner, Communications on pure and applied mathematics **13**, 1 (1960).
 - [36] In his response to Ref. [2] on edge.org, W.D. Willis observes that “Models are interesting precisely because they can take us beyond the data”.
 - [37] Y. Hart, H. Sheftel, J. Hausser, P. Szekely, N. B. Ben-Moshe, Y. Korem, A. Tendler, A. E. Mayo, and U. Alon, Nature methods **12**, 233 (2015).
 - [38] Y. LeCun, Y. Bengio, and G. Hinton, Nature **521**, 436 (2015).
 - [39] Q. V. Le, in *Proceedings of 2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (2013) pp. 8595–8598.
 - [40] J. Tübiana and R. Monasson, Physical Review Letters **118**, 138301 (2017).
 - [41] O. T. Courtney and G. Bianconi, Physical Review E **93**, 062311 (2016).